

Compressing molecular dynamics trajectories: breaking the one-bit-per-sample barrier

Jan Huwald*, Stephan Richter*, Peter Dittrich*

January 13, 2016

Molecular dynamics simulations yield large amounts of trajectory data. For their durable storage and accessibility an efficient compression algorithm is paramount. State of the art domain-specific algorithms combine quantization, HUFFMAN encoding and occasionally domain knowledge.

We propose the high resolution trajectory compression scheme (HRTC) that relies on piecewise linear functions to approximate quantized trajectories. By splitting the error budget between quantization and approximation, our approach beats the current state of the art by several orders of magnitude given the same error tolerance. It allows storing samples at far less than one bit per sample. It is simple and fast enough to be integrated into the inner simulation loop, store every time step, and become the primary representation of trajectory data.

Introduction

Molecular dynamic (MD) simulations are among the largest supercomputer uses. Computing power increases exponentially faster than communication bandwidth.⁶ To retain the ability to durably store, share and even analyze the generated particle trajectories, they have to be represented efficiently.

For example, a recent atomistic model of the SGLT membrane protein, consisting of 90.000 particles simulated for 2.4×10^8 steps (480 ns) generates 259 TiB of raw trajectory data.² The *de-facto* standard approach to handle such large datasets at all, is to down-sample the time-domain of the trajectory to a tiny fraction—in said example by saving

*Biosystem Analysis Research Group

Department of Mathematics and Computer Science and Jena Centre for Bioinformatics

Friedrich Schiller University Jena

Ernst-Abbe-Platz 1-3 / 07743 Jena / Germany

Correspondence: jh@sotun.de / peter.dittrich@uni-jena.de

only $1/50000$ of the steps, except for a few spotlight situations where $1/500$ of all timesteps are saved.

Typically, the down-sampled trajectories are then further compressed. In principle, this is possible using a general purpose lossless compression algorithm, e.g. BZIP2. Unfortunately, general purpose compression suffers from incompressible noise in the less significant bits of the particle positions. They can at most be considered as base line to compare better algorithms against.

MD trajectories are highly amenable to special-purpose compression: The interframe variation of the particle positions is orders of magnitudes smaller than the positions themselves ($\Delta x \ll x$). The demanded precision is typically much smaller than the precision offered by uncompressed representation (32 b or 64 b in IEEE 754). Positions as well as velocities are strongly correlated with values from the past and neighboring particles.

Furthermore, MD trajectory compression has special requirements not fulfilled by general-purpose algorithms. For details, see MARAIS et al.¹⁰ and especially SPÅNGBERG et al.¹⁶ We concentrate on three aspects:

- *Speed*: (de-)compression overhead has to be insignificant compared to the simulation itself to be of any use. The simulator itself is a space-optimal compressor, requiring only the initial state and the elapsed time to be stored. Corollary it has to be *parallelizable*, computable in a *streaming* fashion, have a *small memory footprint*, and elide random writes to the underlying storage.
- *Tunability*: The tolerable error of lossy compression is highly dependent on the simulated scenario and intended analysis. A tunable precision and good performance across all values is thus required.
- *Simplicity*: Complex code fits neither into processor caches nor into programmer minds. It is thus prone to be slow, faulty and not widely implemented.

With the advent of large datasets, a number of compression schemes have been proposed. Most of them are a combination of the following building blocks:

- *Quantization*: the lossy reduction to represent floating point numbers as small integers ($x \mapsto \lfloor 2x/\epsilon_q + 0.5 \rfloor$). The quantization error ϵ_q is a paramount tunable of these algorithms.
- *Delta-coding*: storing the difference of consecutive values instead of the values itself ($(x_0, x_1, x_2, \dots) \mapsto (x_0, x_1 - x_0, x_2 - x_1, \dots)$).
- *Reordering particle coordinates*, so that their consecutive differences are likely to be small. This comes often at a loss of the particle identities: particles of the same element become indistinguishable.
- *Variable-length integer encoding* to store the frequent small integers with few bits without sacrificing the possibility to store rare large values. This is typically achieved using 0^{th} order encoders, which are faster but rely on a fixed distribution of values (e.g., rice coding), or 1^{st} order encoders, which adapt to the observed

distribution to increase the compression rate at the price of higher space and time complexity (e.g., HUFFMAN coding).

All of these techniques have in common that they try to spend less bits per data-point, but keep the number of data-points constant. They thus fail to achieve less than one bit per sample.

SPÅNGBERG et al. proposes TNG-MF1¹⁶, a class of algorithms that use quantization, delta-coding within and between frames, a custom 0th order variable length integer compression, and optionally a combination of BURROW-WHEELER transformation⁴, LEMPEL-ZIV coding¹⁸, and HUFFMAN coding.⁸ MARAIS et al.¹⁰ use quantization, an arithmetic encoder, and interframe prediction with polynomials of order zero or one. Additionally they use *a priori* knowledge about the spatial structure of water to exploit redundancy in adjacent water molecules position and orientation. The venerable XTC file format uses quantization, delta-coding between frames, reordering of coordinates to improve compression of water molecules, and a custom variable-length integer encoding.¹

A completely different approach is realized with the Essential Dynamics tool, which stores particle motions relative to a reference structure in a matrix. The matrix' eigenvectors with the largest eigenvalues are used as a compressed base: a weighted sum of them represents each frame.¹¹ OHTANI et al. represent trajectories by polynomial functions.¹³ A time window for consecutive frames is decreased until a polynomial function fits the data within the given error. Both methods do not support streaming operation and have large time and space overhead. For more examples, see MARAIS ET AL.¹⁰

A related field to MD trajectory compression is the efficient storage of space curves used for geoinformation systems. These algorithms do not store time information, work offline on the entire dataset, and are allowed superlinear runtimes. The prime example is the DOUGLAS-PEUCKER algorithm that iteratively removes points from a curve, as long as they lie within an error corridor between their neighbors.⁵ BELLMAN's algorithm even finds the optimal (minimal error) cover of n points with k lines, but requires $O(n^2)$ time to do so.³

The high resolution trajectory compression algorithm (HRTC) presented here follows a similar, yet faster and simpler approach: Akin to delta-coding, piecewise linear functions are employed to represent trajectories. The resulting support vectors are quantized with a tunable precision and stored using state-of-the-art variable length integer representation. Besides having the highest compression rates and performance, our main novelty is the distribution of the error budget between the quantization and the approximation by functions. Established approaches allow either quantization error^{1,10,16}, or approximation error¹³, but not both at the same time.

Comparing our algorithm to the state of the art is difficult: Crucial simulation parameters are not well documented and a common standard for storing trajectory data needs to be established.⁷ For this purpose we heavily rely on the container-based file format proposed by LUNDBORG et al.:⁹ It allows trajectory storage alongside with parameter values and arbitrary metadata in a single file. It enables usage of different compression algorithms, initially equipped with the option to use XTC, TNG and BZIP2. Both the file format and the code are open source and open to extension. We followed this invitation

and modified the TNG library to support our HRTC compression scheme.

Methodology

The HRTC algorithm we present here uses piecewise linear functions to represent trajectories. We consider only particle positions—noting that for most applications the slope of the approximating function can be used as velocity. Given n particles in a d -dimensional space, we consider each dimension $k \in \{1, \dots, nd\}$ of a state independently. For each dimension k , our algorithm approximates the T points $(x_{k,t})$ of a trajectory by $J(k)$ functions $f_{k,j} : \{1, \dots, \Delta t_{k,j}\} \rightarrow \{\frac{1}{2}z\epsilon_q : z \in \mathbb{Z}\}$ where ϵ_q is the quantization error, and $\Delta t_{k,j}$ is the number of timesteps the function $f_{k,j}$ is covering, with $j \in \{1, \dots, J\}$, $t \in \{1, \dots, T\}$. With a given (maximal) approximation error ϵ_f , our approximation scheme becomes:

$$x_{k,1}, \dots, x_{k,T} \rightsquigarrow f_{k,1}, \dots, f_{k,J(k)} , \quad (1)$$

$$|x_{k,t} - f_{k,j}(t - \sum_{l < j} \Delta t_{k,l})| \leq \epsilon_f , \quad (2)$$

$$\sum_{l < j} \Delta t_{k,l} < t \leq \sum_{l < j+1} \Delta t_{k,l} . \quad (3)$$

A high compression rate can be achieved by covering large durations (maximize $\Delta t_{i,j}$) with functions that require few bits to encode. In contrast to TNG, XTC and other classical algorithms we introduce multiple points of information loss during compression. That is, we split the total error budget ϵ into two parts: the inevitable quantization error ϵ_q , and the approximation error ϵ_f . For this work we use by default $\epsilon_q = \epsilon_f = \frac{1}{2}\epsilon$.

Exploiting the approximation error ϵ_f is our primary vehicle for high compression rates: With $\epsilon_f = 0$ a polynomial representing n points requires n support vectors. Storing them would be no more efficient than storing the points themselves. By introducing an error $\epsilon_f > 0$, multiple functions are valid representations of the point sequence. From this set of functions, we can select those with the fewest support vectors.

In HRTC, we implicitly maintain a set of functions that are valid approximations for each dimension. Every time a new point is added to the sequence, we remove those functions from the set that are not valid approximations of the extended sequence. We continue this process until just before adding the next point would render the candidate set empty. This implements an abstract greedy search for valid approximations covering the maximal timespan $\Delta t_{i,j}$ (see Algorithm 1).

For general functions this algorithm is expensive in time and space: Even if restricted to polynomials with integer coefficients, explicitly storing all functions has exponential space complexity. Generating the new function candidate set F' in line 5 requires at least looking at all Δt values, lifting the lower time bound to $\Omega(n \cdot \Delta t)$.

To remain in our time budget we restrict the function space to linear functions. The first support vector is the final point of the previous interval, the second one at the rightmost point of the current interval. Linear functions allow us to store the entire

Algorithm 1 Abstract algorithm using arbitrary functions for approximation

```

1:  $F \leftarrow$  set of all approximation functions
2:  $\Delta t \leftarrow 0$ 
3: for  $t \in \{0, \dots, T\}$  do
4:    $p' \leftarrow p$  appended by  $\lfloor x_t / (2\epsilon_q) + 0.5 \rfloor$  ▷ Quantize input
5:    $F' \leftarrow$  valid approximation functions for  $p'$ 
6:   if  $F' = \emptyset$  or  $t = T$  then
7:     output  $(F, \Delta t)$ 
8:      $p \leftarrow (x_t)$ 
9:      $F \leftarrow$  set of all functions
10:     $\Delta t \leftarrow 0$ 
11:   else
12:      $p \leftarrow p'$ 
13:      $F \leftarrow F'$ 
14:      $\Delta t \leftarrow \Delta t + 1$ 
15:   end if
16: end for

```

candidate set F using two integers, and merging as well as computing the candidate set in constant time.

The set of linear functions through a point x that are valid approximations for a second point x' separated by Δt timesteps form a $2\epsilon_f$ wide error cone around x' (see Figure 1). The set is denoted by AF:

$$\begin{aligned}
& \{f : f(0) = x \wedge |f(\Delta t) - x'| \leq \epsilon_f\} \\
& \supseteq \{t \mapsto \frac{t}{\Delta t}v + x : v \in [v_\perp, v_\top]\} \\
& =: \text{AF}(v_\perp, v_\top) \quad \text{with } v_\perp = \frac{x' - x - \epsilon_f}{\Delta t}, \quad v_\top = \frac{x' - x + \epsilon_f}{\Delta t}.
\end{aligned} \tag{4}$$

The set of functions is completely represented by x , v_\perp , and v_\top . The intersection of multiple such sets that share a common support vector x but differ in extremal slopes $v_{\perp,i}$ and $v_{\top,i}$ can be merged efficiently:

$$\bigcap_i \text{AF}(x, v_{\perp,i}, v_{\top,i}) = \text{AF}(x, \max_i v_{\perp,i}, \min_i v_{\top,i}). \tag{5}$$

This intersection contains all valid approximations for *all* input points. Note that the intersection is empty if the lower bound v_\perp becomes larger than the upper bound v_\top , that is, $\text{AF}(x, v_\perp, v_\top) = \emptyset$ iff. $v_\perp > v_\top$. The above formulation allows to incrementally update the function candidate set. The sequence of points of the current interval does not have to be stored anymore.




Figure 1: Illustration of a segment used to capture a part of the trajectory: The left-most node x is the left support vector of the current one. Together with the error bars around each node $x' \dots x'''$, it induces the extremal slopes $v_{\perp,i}, v_{\top,i}$ (dotted lines). The shaded cones cover all possible slopes for each point. While successively including more points into the segment, the range of valid slopes $[v_{\perp}, v_{\top}]$ (red line) decreases. The error cone induced by x''' (shaded red) does not intersect with $[v_{\perp}, v_{\top}]$ covering x, x' and x'' (shaded dark-grey). Thus $v_{\perp} \Delta t''$ becomes the terminal node of this segment: it is the valid point closest to x'' .

Combined storage of multiple trajectories

A single trajectory can be stored in memory as a sequence $(v_0, \Delta t_0), (v_1, \Delta t_1), \dots$. When storing multiple trajectories, we have to counter the issue that their support vectors are not synchronized in time. A naive approach would include a trajectory index k_i for each support vector: $(k_1, v_{k_1}, \Delta t_{k_1}), (k_2, v_{k_2}, \Delta t_{k_2}), \dots$. However, this would require $O(\log nd)$ additional space per support vector, which we avoid with the following procedure.

We exploit the fact that the start time and duration of a segment imply the time at which the next support vector has to be expected. To use this insight, we maintain an auxiliary priority queue Q_{expected} , which stores tuples (t, k) sorted by time t and secondarily trajectory index k . The minimal element of the queue denotes the support vector to be outputted next. The queue always contains nd elements, thus the space gain is traded against an additional time complexity of $O(\log nd)$ per inserted interval.

To reorder the support vectors from the sequence of discovery to the sequence of storage described above, we use a second priority queue Q_{known} . It stores all support vectors that can not be stored immediately because a preceding support vector (according to Q_{expected}) is not yet known. It contains up to $O(nd \max_i \Delta t_i)$ elements. This implies a memory overhead of the same magnitude in the compressor. The time overhead can be reduced to $O(nd)$ assuming time points are dense—if the number of support vectors of all trajectories exceeds the number of time points. Due to the blockwise compression and limited Δt observed, this overhead is much less dramatic in practice.

A key frame is used to initialize the queue and to provide the initial support vectors

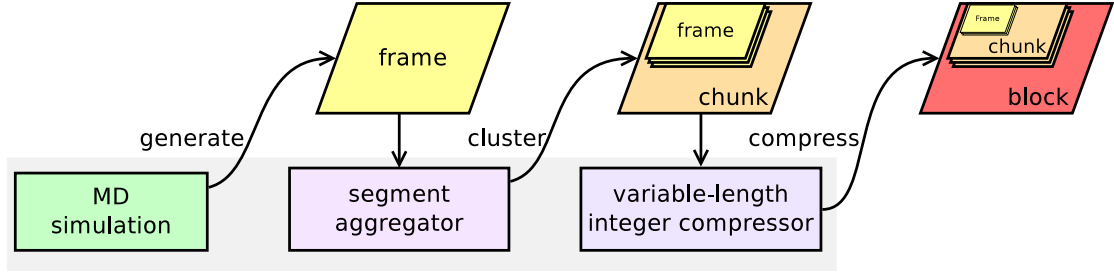


Figure 2: Hierarchical storage of MD simulation generated data. Top row denotes data, bottom row code. Each block starts with a key frame. A high-performance implementation typically integrates all lower boxes.

for each dimension. The key frame stores the quantized values of $x_{0,0}, x_{1,0}, \dots$ with a fixed bit count, without further compression. The number of bits $\lceil \log \frac{L}{2\epsilon_q} \rceil$ is determined by the quantization granularity ϵ_q and the edge length L of the bounding box of the trajectory.

Hierarchy of computation and storage

The HRTC algorithm has to conciliate between two opposing forces: Memory pressure and data cache locality demand to minimize the window of trajectory data held in memory during compression. Higher compression rates require a larger timespan of data to reason about. To minimize cache thrashing, switching between simulation and compression should happen at a low frequency—asking for a large window, too. At last, to generate a searchable compressed data-stream, key frames have to be inserted at regular time intervals.

Balancing these demands has led to the following hierarchy of computation and storage (cf. Figure 2): An MD simulator calls the compression library for each generated timestep—denoted a *frame*. The data is approximated with linear functions, their resulting support vectors are buffered. Once a threshold buffer size is reached (e.g., 8 KB storing 1024 support vectors), all support vectors are fed through the variable length integer encoding, yielding one *chunk*. The largest unit—a *block*—contains a user-specified number of frames. It starts with a single key frame and contains all chunks belonging to the encoded frames. The block size is tuned by the user depending on the desired compression and seek time.

The HRTC compression integrated into the TNG library does not fully exploit this hierarchy. Due to design constraints of the TNG library, all frames of a block are collected and then compressed at once. When speed matters, HRTC could be used directly.

The update compression algorithm, together with storage hierarchy and multiple trajectory storage are described by Algorithm 2. The respective decompression is described by Algorithm 3.

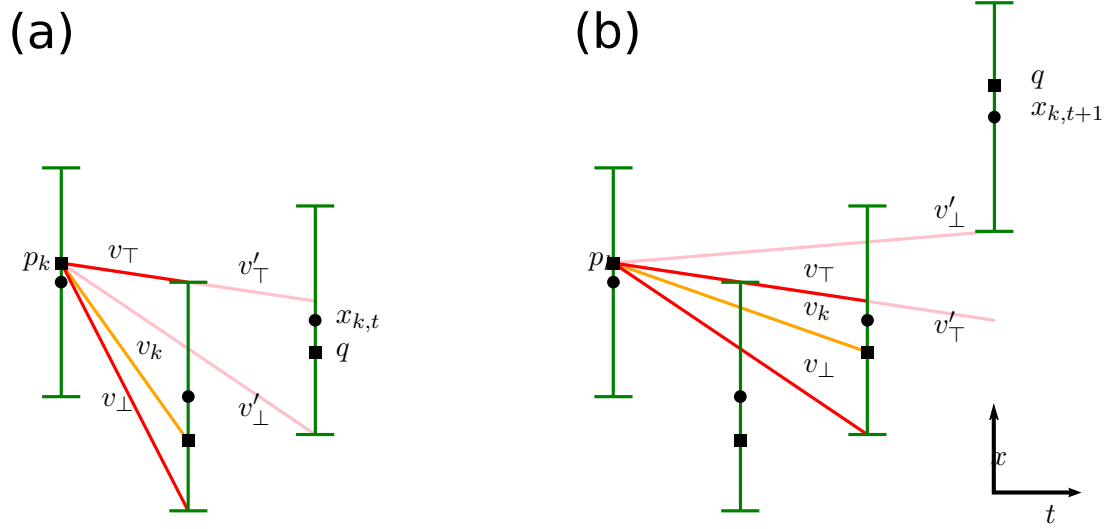


Figure 3: Illustration of the variables used during compression for each dimension k as used in algorithm 2. Shown are two cases: A point to be added either extends the segment (a), or causes termination of the segment at the previous point (b). The input $x_{k,t}$ is shown as circle, the quantized values as square. Four variables are stored per dimension: p_k denotes the quantized support vector starting the segment, v_k the slope to the last point added, and v_{\perp} and v_{\top} the lower and upper slope bounds. In addition, three temporary variables are used: the quantized input value q , and the updated slope bounds v'_{\perp} and v'_{\top} .

Algorithm 2 The HRTC compression algorithm. For an illustration of the variables used, see Figure 3.

```

1: Input: Trajectory  $(x_{t,k})$  for  $t \in \{0, \dots, T\}, k \in \{1, \dots, nd\}$ , quantization error  $\epsilon_q$ ,
   approximation error  $\epsilon_f$ 
2: Output: Compressed trajectory (incl. key frames)
3: global  $v_{\perp,k}, v_{\top,k}, v_k, p_k, \Delta t_k$  for  $k \in \{1, \dots, nd\}$   $\triangleright$  current approximation
4: global  $q$   $\triangleright$  current quantized input  $x_{t,k}$ 
5: global outputBuf = ()
6: global  $Q_{\text{known}} \leftarrow \emptyset, Q_{\text{expected}} \leftarrow \emptyset$ 

7: for  $t \in \{0, \dots, T\}$  do  $\triangleright$  For each input frame
8:   DRAINQUEUE  $\triangleright$  Move expected known segments to output buffer
9:   if  $t \bmod \text{blockSize} = 0$  then  $\triangleright$  ENCODE KEY FRAME
10:    while  $Q_{\text{expected}} \neq \emptyset$  do  $\triangleright$  While there are segments to end
11:       $(t', k) \leftarrow \min(Q_{\text{expected}})$   $\triangleright$  Get time and dimension of next segment
12:      FLUSHSEGMENT( $t', k$ )  $\triangleright$  End this segment and push it into  $Q_{\text{known}}$ 
13:    end while
14:    DRAINQUEUE  $\triangleright$  Move  $Q_{\text{known}}$  (known segments) to output buffer
15:    FLUSHCHUNK  $\triangleright$  Compress output buffer and output it
16:     $Q_{\text{expected}} \leftarrow \{(t, 0), \dots, (t, nd - 1)\}$   $\triangleright$  Initialize  $Q_{\text{expected}}$  with current time
17:    for  $k \in \{1, \dots, nd\}$  do  $\triangleright$  For each dimension  $k$ :
18:       $p_k \leftarrow \lfloor x_{t,k} / (2\epsilon_q) + 0.5 \rfloor$   $\triangleright$  get the next input, quantize and remember it
    as  $p_k$ ,
19:      output  $p_k$   $\triangleright$  and output it (for the key frame).
20:       $v_{\perp,k} = -\infty, v_{\top,k} = \infty$   $\triangleright$  initialize empty segment
21:       $\Delta t_k \leftarrow 0$ 
22:    end for
23:  else  $\triangleright$  ENCODE SEGMENTS
24:    for  $k \in \{1, \dots, nd\}$  do  $\triangleright$  for each dimension  $k$ 
25:       $q \leftarrow \lfloor x_{t,k} / (2\epsilon_q) + 0.5 \rfloor$   $\triangleright$  quantize next input
26:       $v'_{\top} \leftarrow \min(v_{\top,k}, (q - p_k + \epsilon_f) / \Delta t_k)$   $\triangleright$  Computer lower and upper bound
27:       $v'_{\perp} \leftarrow \max(v_{\perp,k}, (q - p_k - \epsilon_f) / \Delta t_k)$   $\triangleright$  of the approximating functions.
28:      if  $v'_{\top} < v'_{\perp}$  then  $\triangleright$  If no valid approximating function remains
29:        FLUSHSEGMENT( $t, k$ )  $\triangleright$  terminate current segment with previous
        point.
30:      else
31:         $v_{\perp,k} \leftarrow v'_{\perp}, v_{\top,k} \leftarrow v'_{\top}$   $\triangleright$  Update current approximation
32:         $v_k \leftarrow (q - p_k) / \Delta t_k$ 
33:         $\Delta t_k \leftarrow \Delta t_k + 1$ 
34:      end if
35:    end for
36:  end if
37: end for

```

```

38: function FLUSHSEGMENT( $t, k$ )
39:    $\Delta q \leftarrow \begin{cases} \Delta t_k v_{\perp, k} & \text{if } v_k < v_{\perp, k} \\ \Delta t_k v_k & \text{if } v_k \in [v_{\perp, k}, v_{\top, k}] \\ \Delta t_k v_{\top, k} & \text{if } v_k > v_{\top, k} \end{cases}$   $\triangleright$  Choose support vector closest to prev.
      point.
40:   insert  $(t - \Delta t_k, k, \Delta t_k, \Delta q)$  into  $Q_{\text{known}}$   $\triangleright$  Insert support vector to  $Q_{\text{known}}$ ,
       $\triangleright$  which is sorted by  $(t - \Delta t_k, k)$ .
41:    $\Delta t_k \leftarrow 1$   $\triangleright$  Initialize next segment duration with 1
42:    $v_{\perp, k} \leftarrow q - p_k - \epsilon_f$   $\triangleright$  Reset upper and lower bounds
43:    $v_{\top, k} \leftarrow q - p_k + \epsilon_f$ 
44:    $p_k \leftarrow p_k + \Delta q$   $\triangleright$  Set starting point of next segment
       $\triangleright$  to terminal point of current segment.
45: end function

46: function FLUSHCHUNK
47:   compress outputBuf using VSE-R, output result
48:   outputBuf  $\leftarrow ()$ 
49: end function

50: function DRAINQUEUE( $Q_{\text{known}}, Q_{\text{expected}}, \text{outputBuf}$ )
51:   while  $\min(Q_{\text{expected}}) \stackrel{t, k}{=} \min(Q_{\text{known}})$  do  $\triangleright$  While next segment is known
52:      $(t, k, \Delta t, \Delta q) \leftarrow \text{extractMin}(Q_{\text{known}})$   $\triangleright$  Get next known segment
53:      $\text{extractMin}(Q_{\text{expected}})$   $\triangleright$  Update the expected start time
54:     insert  $(t + \Delta t, k)$  into  $Q_{\text{expected}}$ 
55:     outputBuf  $\leftarrow (\Delta q, \text{outputBuf}, \Delta t)$   $\triangleright$  Append segment to output buffer
       $\triangleright$  (see section “chunk encoding”)
56:     if length of outputBuf  $\geq 2 \cdot \text{chunkLength}$  then  $\triangleright$  If chunk is full.
57:       FLUSHCHUNK  $\triangleright$  compress and output chunk.
58:     end if
59:   end while
60: end function

```

Algorithm 3 The HRTC **d**ecompression algorithm. The handling of chunks for integer compression has been omitted for brevity—it is implied in line 11.

```

1: Input: Compressed trajectory (incl. key frames)
2: Output: Uncompressed trajectory  $(x_{t,k})$ 
3: for  $t \in \{0, \dots, T\}$  do
4:   if  $t \bmod \text{blockSize} = 0$  then ▷ decode key frame
5:     for  $k \in \{1, \dots, nd\}$  do
6:       read  $q$ 
7:        $p_k \leftarrow q\epsilon_q, t_k \leftarrow t, v_k \leftarrow 0, \Delta t_k \leftarrow 0$  ▷ Set starting point of segment
8:     end for
9:      $Q_{\text{expected}} \leftarrow \{(t, 1), \dots, (t, nd)\}$  ▷ Initialize  $Q_{\text{expected}}$  with current time
10:  end if
11:  while  $t \stackrel{t}{=} \min(Q_{\text{expected}})$  do ▷ For all segments starting at  $t$ 
12:     $(t, k) \leftarrow \text{extractMin}(Q_{\text{expected}})$  ▷ Get dimension of next segment
13:    read  $(d, q)$  using VSE-R decompression ▷ Get its duration and support
    vector
14:     $p_k \leftarrow p_k + \Delta t_k v_k$  ▷ Compute starting point
15:     $v_k \leftarrow q\epsilon_q/d$  ▷ Compute slope of segment
16:     $t_k \leftarrow t$  ▷ Update time of segment
17:     $\Delta t_k \leftarrow d$  ▷ Duration of segment
18:    insert  $(t + d, k)$  into  $Q_{\text{expected}}$  ▷ Update expected time for next segment
19:  end while
20:  for  $k \in \{1, \dots, nd\}$  do ▷ Interpolate current frame
21:    output  $x_{t,k} \leftarrow p_k + (t - t_k)v_k$  ▷ and output it.
22:  end for
23: end for

```

Integer compression and chunk encoding

For variable-length integer encoding, we use the `INTEGER ENCODING LIBRARY`.¹⁷ It offers several codecs. After selecting for time and speed, we chose the codec `VSE-R`.¹⁵ The library only encodes unsigned integers. Where signed integers occur in our algorithm they are mapped to unsigneds:

$$i \mapsto \begin{cases} 2i & \text{if } i \geq 0, \\ -2i + 1 & \text{if } i < 0. \end{cases} \quad (6)$$

`VSE-R` encodes groups of consecutive integers with the number of bits required by the largest element of the group. The optimal length of the group is computed using dynamic programming. This allows storing the number of bits only once for several integers to be stored. It is the basis for high performance of `VSE-R` (regarding throughput and compression). It also means that the performance is suboptimal when encoding integers of alternating magnitude. In our case the magnitude of space and time deltas can be different by several orders of magnitude. So we rearrange the support vectors such that time and space deltas are grouped together, respectively. Instead of the queue $(\Delta x_1, \Delta t_1), \dots, (\Delta x_n, \Delta t_n)$ we store the doubled ended queue $\Delta x_n, \dots, \Delta x_1, \Delta t_1, \dots, \Delta t_n$ (see function `DRAINQUEUE` in Algorithm 2).

Optional adaptation for deep simulator integration

Because the computational demands for our compression method are small, we envisage future integration of it into the inner loop of MD simulation programs: Every update of a particle’s state is immediately followed by the compression of the new position. `HRTC` can then serve as the primary mechanism to retrieve simulation data. In the following, we describe optional adaptations of our algorithm for this purpose, which are, however, not applied for our performance evaluation in the result section.

For a state-of-the-art MD application the integration of `HRTC` in its inner loop poses additional challenges. High performance MD simulators rely on specialized hardware—from GPUs to custom ASICs.¹⁴ The performance characteristic of these platforms differs from a typical desktop CPU. Non-uniform memory access, diverging control flow and branches are much more expensive compared to arithmetic operations. Addition and multiplication are especially fast compared to division and other mathematical operations.¹² To accommodate `HRTC` compression on these machines, we can adapt the algorithm. The critical section of our algorithm is the check whether the current candidate set of curves is empty after adding the next point. We reformulate it to avoid division operations, and rely on conditional writes instead of branches. Then the special purpose hardware only needs to transfer id and position of those dimensions where the check failed. The host computer executes all further compression steps (sorting, queue management, variable-length integer encoding) in parallel to the kernel running on special purpose hardware.

The original condition whether a point does not lie in the current set of linear functions

is:

$$v'_\perp > v'_\top, \quad (7)$$

$$v'_\perp = \max(v_\perp, \frac{x' - x - \epsilon_f}{\Delta t}), \quad (8)$$

$$v'_\top = \min(v_\top, \frac{x' - x + \epsilon_f}{\Delta t}). \quad (9)$$

This code requires 4 additions, 2 divisions, 3 comparisons and 3 branches. Equation 7-9 can be merged and simplified by case analysis:

$$v'_\perp > v'_\top \Leftrightarrow \begin{cases} v_\perp > \frac{x' - x + \epsilon_f}{\Delta t} & \text{if } v_\perp \geq \frac{x' - x - \epsilon_f}{\Delta t} \wedge v_\top \geq \frac{x' - x + \epsilon_f}{\Delta t}, \quad (10.1) \\ v_\perp > v_\top & \text{if } v_\perp \geq \frac{x' - x - \epsilon_f}{\Delta t} \wedge v_\top < \frac{x' - x + \epsilon_f}{\Delta t}, \quad (10.2) \\ \frac{x' - x - \epsilon_f}{\Delta t} > \frac{x' - x + \epsilon_f}{\Delta t} & \text{if } v_\perp < \frac{x' - x - \epsilon_f}{\Delta t} \wedge v_\top \geq \frac{x' - x + \epsilon_f}{\Delta t}, \quad (10.3) \\ \frac{x' - x - \epsilon_f}{\Delta t} > v_\top & \text{if } v_\perp < \frac{x' - x - \epsilon_f}{\Delta t} \wedge v_\top < \frac{x' - x + \epsilon_f}{\Delta t}, \quad (10.4) \\ \text{false} & \text{otherwise.} \end{cases} \quad (10)$$

Case 10.2 is impossible: it implies that the curve set was already empty after insertion of the previous point. That would already have been remedied by starting a new segment. Case 10.3 is impossible as all variables are strictly positive. This allows us to rewrite equation 10 without expensive division operations: Instead of the extremal slopes v_\perp, v_\top we store the time $\Delta t_\perp, \Delta t_\top$ and value x_\perp, x_\top of the previous extrema. Then a new segment starts iff.

$$\begin{aligned} & \frac{x_\perp - x + \epsilon_f}{\Delta t_\perp} > \frac{x' - x + \epsilon_f}{\Delta t} \bigwedge \\ & \frac{x_\top - x - \epsilon_f}{\Delta t_\top} < \frac{x' - x - \epsilon_f}{\Delta t} \\ \Leftrightarrow & (x_\perp - x + \epsilon_f)\Delta t > (x' - x + \epsilon_f)\Delta t_\perp \bigwedge \\ & (x_\top - x - \epsilon_f)\Delta t < (x' - x - \epsilon_f)\Delta t_\top. \end{aligned} \quad (11)$$

The update of the stored extrema can be modified in the same way. Update $x_\perp \leftarrow x'$ and $t_\perp \leftarrow t_1$ when

$$\begin{aligned} & v_\perp < \frac{x' - x - \epsilon_f}{\Delta t} \\ \Leftrightarrow & \frac{x_\perp - x' - \epsilon_f}{\Delta t_\perp} < \frac{x' - x - \epsilon_f}{\Delta t} \\ \Leftrightarrow & (x_\perp - x' - \epsilon_f)\Delta t < (x' - x - \epsilon_f)\Delta t_\perp. \end{aligned} \quad (12)$$

Analogous for x_\top . Implementing eq. 11 and 12 requires 13 additions, 6 multiplications, 2 conditional moves and one branch. Although the operation count is higher, expensive division and branch operations have been omitted.

Results and discussion

A library implementing the HRTC algorithm as described above is available under a GPL-3 open source license at <https://github.com/biosystemanalysis/hrtc>. To compare our compression algorithm with the state of the art, we additionally integrated it into the TRAJECTORY-NG library. The merged library is available at <https://github.com/biosystemanalysis/tng>.

Hrtc outperforms existing compression methods

TNG comes with a benchmark application, used to compare compression algorithms, here. The benchmark applies velocity verlet integration to simulate 512 particles with a harmonic well potential $U(\Delta x) = \sin(\min(\|\Delta x\|, \frac{\pi}{2}))^2$ and an arbitrary mass 2, and timestep $2 \cdot 10^{-4}$ with arbitrary units (a.u.). Initially, the particles are distributed randomly in a $15 \times 16 \times 17$ (a.u.) cuboid and equilibrated for 10^5 steps. After equilibration, this simulation is run for 10^7 more timesteps to generate the benchmark trajectory. The transient phase is omitted in order to avoid artifacts when applying different sub-sampling rates. The resulting benchmark trajectory is then compressed using either TNGs native compression algorithm or our HRTC compression. We compare compression rates while varying the sub-sampling rate, the number of frames per block, and the maximal error ϵ . The results are depicted in Table 1. The remaining parameters are held constant during this paper: chunk size is set to 1024 support vectors, and block size to 2048 frames.

HRTC outperforms the TNG compression in all cases, except when a very coarse sub-sampling rate ($\lesssim 1 : 1024$) is combined with high spatial resolution ($\epsilon \lesssim 0.001$). These cases are practically irrelevant for two reasons: First, in our example $\epsilon = 0.0001$ corresponds to 18 b of position information per dimension—almost equal to the 24 b of a single precision float mantissa. And second, the position inaccuracy introduced by low temporal resolution far exceeds the error bound 0.0001 even for minuscule particle velocities.

Furthermore, HRTC appears to outperform all compression methods investigated by MARAIS, ET AL.¹⁰ As neither their implementation nor their test data is public, we have to cautiously compare our benchmarks despite them running on different datasets. The best compression rate MARAIS’ algorithms achieves is 20.8 with 1:2 sub-sampling, a 12 b quantization and 0.014 Å positional error. This rate is overachieved by HRTC already at the much coarser 1:32 sub-sampling with much smaller error bound $\epsilon = 0.0001$ —yielding a large buffer to compensate for the different datasets being used. At a comparable error rate ($\epsilon = 0.01$) and sub-sampling (1:2) HRTC achieves a compression ratio of 3419.

Compressing below one bit per sample

To demonstrate the ability of our algorithm to reduce storage space needed per sample below the 1 b limit, we compress a 2048 frame 512 particle simulation with varying error rates ϵ . For values of $\epsilon \gtrsim 0.0000006$ the average space demanded per sample drops below

Table 1: Comparison of compression ratio (uncompressed size / compressed size) for different error bounds ϵ and sub-sampling rates. The size of the original, uncompressed TNG file was 114 GiB. The compressed file sizes vary between 796 KiB and 7.0 GiB. Note that the sub-sampling factor is also included in the compression ratio.

Sampling rate	$\epsilon = 1$		$\epsilon = 0.1$		$\epsilon = 0.01$		$\epsilon = 0.001$		$\epsilon = 0.0001$	
	HRTC	TNG	HRTC	TNG	HRTC	TNG	HRTC	TNG	HRTC	TNG
1:1	4709	38	3644	38	2507	38	1108	33	411	16
1:2	9045	77	6578	76	3419	74	1326	60	468	22
1:4	16 618	153	10 207	152	4195	144	1506	99	531	38
1:8	28 410	307	13 490	303	4904	273	1711	150	606	66
1:16	42 176	612	16 338	598	5597	499	1940	204	699	116
1:32	54 303	1225	18 423	1171	6358	854	2232	308	782	209
1:64	64 719	2438	20 796	2239	7193	1325	2546	537	845	379
1:128	73 421	4825	23 947	4146	8203	1844	2828	945	915	692
1:256	82 726	9458	27 176	7287	9300	2621	3078	1680	1311	1205
1:512	99 107	19 225	32 324	12 222	11 020	4757	3791	3182	2351	2370
1:1024	128 881	40 580	41 306	19 808	13 915	9294	6401	6466	4672	5257
1:2048	150 901	72 887	48 670	26 458	17 832	16 985	10 604	12 633	8348	9862




Figure 4: Average number of bits required to compress one sample—the position of one particle at one timepoint—depending on error bound ϵ . A trajectory of 512 particles over 10^7 timesteps generated by the TNG benchmark application was compressed with 2048 frames per block.

Table 2: Runtime (s) of an MD simulation with the trajectory stored either with TNG, HRTC, or not at all. The storage precision ϵ is varied, but has almost no effect on the runtime.

trajectory storage	$\epsilon = 0.0001$	$\epsilon = 0.001$	$\epsilon = 0.01$	$\epsilon = 0.1$	$\epsilon = 1$
no storage	1727.8	1730.1	1729.3	1727.0	1728.8
HRTC compression	1726.9	1727.2	1723.5	1728.4	1723.7
TNG compression	1735.5	1723.5	1732.5	1729.9	1726.6

1 b (see Figure 4).

Compression is fast

We measured the throughput of compression and decompression with the same benchmark simulation used for the compression rate estimation. As this simulation is executed on a general-purpose CPU, it requires relatively large amounts of time to compute the pairwise forces in each step. This made the overhead of compression statistically insignificant compared to the run-time variations of the simulation itself (see Table 2). To compare the throughput of different algorithms, we thus measured the time taken to compress a 114 GiB trajectory that was stored uncompressed on an SSD.

The throughput of the tested algorithms depends on the magnitude of change of the trajectory to compress. To simulate trajectories with different speed, we sub-sampled our test trajectory. A $1:n$ subsampling results in a n -fold speedup of the trajectory to compress. The throughput of TNG and HRTC compression are depicted in Figure 5. For slowly varying trajectories (speed-up below 100), HRTC performs around 100 MiB/s. For speed-ups below factor 500 HRTC beats TNG in terms of throughput. For fast rates of change, HRTC’s throughput converges against 20 MiB/s. On the same dataset, the general purpose BZIP2 compression achieves no more than 7 MiB/s.

On their own data, MARAIS et al. report compression rates between 13 and 39 MiB/s.¹⁰

A second test with naked HRTC compression—without the overhead of integration into the TNG library—reveals the extremal throughput: A trajectory with constant particle position is compressed with 520 MiB/s and decompressed with 1838 MiB/s. Trajectories with purely random positions reach 65 MiB/s and 40 MiB/s for compression and decompression respectively.

All throughput tests have been performed on a single core of an Intel Xeon E5-2690 2.9 GHz CPU.




Figure 5: Compression throughput in MiB/s of input data processed depending on trajectory velocity. Trajectories with faster movement were obtained by sub-sampling a common trajectory: A speedup of n is equal to taking every n th frame.

Effect of error distribution between ϵ_q and ϵ_f

In the preceding sections we assumed the error budget ϵ is equally distributed between function approximation and quantization ($\epsilon_q = \epsilon_f = \frac{1}{2}\epsilon$). The main motivation for that is to keep the number of parameters as small as possible. To test whether this distribution is suitable, we parametrized the error distribution by $\lambda \in (0, 1]$:

$$\epsilon_q = \lambda\epsilon, \quad \epsilon_f = (1 - \lambda)\epsilon_q \quad (13)$$

We then compressed the TNG test data used above at different sub-sampling rates. We compared the compression ratio achieved while varying λ . The result is depicted in Figure 6. Between $0.01 \leq \lambda \leq 0.5$ a plateau of the compression ratio can be seen independent of the sub-sampling rate. When spending more the 50% of the error budget on quantization the compression rate sharply decreases. This effect is the more pronounced, the more volatile the trajectory is (at sub-sampling rates below 1 : 16). We have observed similar behavior for all other datasets we tested (data not shown).

Conclusions

We have developed a novel compression algorithm specifically for storage of molecular dynamic trajectories. The algorithm is lossy, with a user-specified error bound. By splitting the available error budget between quantisation error and function approximation error, we attain previously unachievable compression rates far below one bit per sample.

Even when saving with high fidelity (small time steps) compression rate and throughput are outstandingly high. Thus, we propose to use our format as primary representation of simulation data coming out of MD simulation kernels. This will reduce the bandwidth demands between simulation kernel and analysis tools. Using the computed linear functions as primal data representation will allow more integrated queries than the currently used uncompressed snapshots: for example a check for the minimal distance between two particles can be answered analytically on the level of segments, instead of iterating over all points.

Our approach offers some simple, yet rewarding extension points. Foremost, parallelization of the compression is trivial: just compress subsets of the particles independently, preferably on the same cores that compute these subsets. Due to better use of caches and smaller amounts of buffered support vectors when waiting for the termination of long-lasting function segments, we expect a super-linear speed-up from parallel execution.

Secondly, performance gains can be achieved by tight integration of function approximation and variable length integer representation (we used an external library so far). This will save a whole pass over the data and allows to tune the integer compression parameters to the expected distribution occurring in our use-case. Alternatively, an adaptive coder (e.g. entropy coder or arithmetic coder) could be employed. While this would result in a modest improvement of the compression rate (we estimate at most factor 2), it would also reduce the throughput of our compression significantly.




Figure 6: Compressed size depending on the error distribution coefficient λ relative to the uncompressed file. Given the quantization error $\epsilon_q = \lambda\epsilon$ and the approximation error $\epsilon_f = (1 - \lambda)\epsilon$, for $\lambda \approx 0$ the algorithm degrades to a kind of linear regression, and for $\lambda \approx 1$ all error budget is used for quantization. The algorithm then degrades to a linear extrapolation scheme similar to the one used by MARAIS ET AL.¹⁰ The uncompressed data contained 10^4 frames of 512 particles with 3 dimensions. It was sub-sampled at eight different rates (1:1 - 1:64). Compression was performed with total error $\epsilon = 0.01$, block size 10^4 using plain HRTC (without TNG metadata).

Thirdly, integrating the TNG library and HRTC more tightly should offer significant performance benefits. We demonstrated that the throughput of our naked HRTC library is much higher than our current integration into the TNG library. There is an impedance mismatch between both libraries. A deeper integration of HRTC into TNG accompanied with some changes of TNG’s architecture to allow calling it inside the inner-most simulation loop without performance drawbacks will lift trajectory compression to a new level.

Author Contributions

Jan Huwald designed the algorithm and implemented it in the HRTC library. Stephan Richter integrated HRTC into TNG, generated the test data and executed the benchmarks. The first draft was written jointly by Jan Huwald and Stephan Richter. Peter Dittrich was responsible for project supervision and intensive revision of the draft.

Acknowledgments

The authors acknowledge support from the European Union through funding under FP7-ICT-2011-8 project HIERATIC (316705).

References

- [1] ABRAHAM, M. J., MURTOLA, T., SCHULZ, R., PÁLL, S., SMITH, J. C., HESS, B., AND LINDAHL, E. Gromacs: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* 1–2 (2015), 19 – 25.
- [2] ADELMAN, J., SHENG, Y., CHOE, S., ABRAMSON, J., WRIGHT, E., ROSENBERG, J., AND GRABE, M. Structural determinants of water permeation through the sodium-galactose transporter vsGLT. *Biophys. J.* 106, 6 (2014), 1280 – 1289.
- [3] BELLMAN, R. On the approximation of curves by line segments using dynamic programming. *Commun. ACM* 4, 6 (June 1961), 284–.
- [4] BURROWS, M., AND WHEELER, D. J. A block-sorting lossless data compression algorithm. Tech. rep., Digital, Systems Research Center, Palo Alto, CA, USA, 1994.
- [5] DOUGLAS, D. H., AND PEUCKER, T. K. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Cartographica: The International Journal for Geographic Information and Geovisualization* 10, 2 (1973), 112–122.
- [6] ESMAEILZADEH, H., BLEM, E., ST. AMANT, R., SANKARALINGAM, K., AND BURGER, D. Dark silicon and the end of multicore scaling. In *ISCA ’11: Proceedings*

of the 38th Annual International Symposium on Computer Architecture (New York, NY, USA, 2011), ACM, pp. 365–376.

- [7] HINSEN, K. Mosaic: A data model and file formats for molecular simulations. *J. Chem. Inf. Model.* *54*, 1 (2014), 131–137.
- [8] HUFFMAN, D. A method for the construction of minimum-redundancy codes. *Proc. IRE* *40*, 9 (Sept 1952), 1098–1101.
- [9] LUNDBORG, M., APOSTOLOV, R., SPÅNGBERG, D., GÄRDENÄS, A., VAN DER SPOEL, D., AND LINDAHL, E. An efficient and extensible format, library, and api for binary trajectory data from molecular simulations. *J. Comput. Chem.* *35*, 3 (2014), 260–269.
- [10] MARAIS, P., KENWOOD, J., SMITH, K. C., KUTTEL, M. M., AND GAIN, J. Efficient compression of molecular dynamics trajectory files. *J. Comput. Chem.* *33*, 27 (2012), 2131–2141.
- [11] MEYER, T., FERRER-COSTA, C., PÉREZ, A., RUEDA, M., BIDON-CHANAL, A., LUQUE, F. J., LAUGHTON, C. A., , AND OROZCO, M. Essential dynamics: a tool for efficient trajectory compression and management. *J. Chem. Theory Comput.* *2*, 2 (2006), 251–258.
- [12] NVIDIA. CUDA C Programming Guide (version 7.5). <http://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>, September accessed November 2, 2015.
- [13] OHTANI, H., HAGITA, K., ITO, A. M., KATO, T., SAITOH, T., AND TAKEDA, T. Irreversible data compression concepts with polynomial fitting in time-order of particle trajectory for visualization of huge particle system. *J. Phys.: Conf. Ser.* *454*, 1 (2013), 012078.
- [14] SHAW, D. E., DENEROFF, M. M., DROR, R. O., KUSKIN, J. S., LARSON, R. H., SALMON, J. K., YOUNG, C., BATSON, B., BOWERS, K. J., CHAO, J. C., EASTWOOD, M. P., GAGLIARDO, J., GROSSMAN, J. P., HO, C. R., IERARDI, D. J., KOLOSSVÁRY, I., KLEPEIS, J. L., LAYMAN, T., MCLEAVEY, C., MORAES, M. A., MUELLER, R., PRIEST, E. C., SHAN, Y., SPENGLER, J., THEOBALD, M., TOWLES, B., AND WANG, S. C. Anton, a special-purpose machine for molecular dynamics simulation. In *Proceedings of the 34th Annual International Symposium on Computer Architecture* (New York, NY, USA, 2007), ISCA '07, ACM, pp. 1–12.
- [15] SILVESTRI, F., AND VENTURINI, R. VSEncoding: efficient coding and fast decoding of integer lists via dynamic programming. In *Proceedings of the 19th ACM international conference on Information and knowledge management* (New York, NY, USA, 2010), ACM, pp. 1219–1228.

- [16] SPÅNGBERG, D., LARSSON, D. S. D., AND VAN DER SPOEL, D. Trajectory NG: portable, compressed, general molecular dynamics trajectories. *J. Mol. Model.* *17*, 10 (2011), 2669–2685.
- [17] YAMAMURO, T., SILVESTRI, F., AND VENTURINI, R. Integer encoding library. Website <http://integerencoding.isti.cnr.it> (accessed Jan 7, 2015) with source code at <https://github.com/maropu/integer-encoding-library>, 2011-2012.
- [18] ZIV, J., AND LEMPEL, A. A universal algorithm for sequential data compression. *IEEE Trans. Inf. Theory* *23*, 3 (May 1977), 337–343.